

A Randomized Rounding Approach to the Traveling Salesman Problem

Shayan Oveis Gharan *

Management Science and Engineering Management Science and Engineering
Stanford University
Stanford, CA 94305
shayan@stanford.edu

Amin Saberi +

Stanford University
Stanford, CA 94305
saberi@stanford.edu

Mohit Singh

School of Computer Science
McGill University
Montreal, QC, Canada, H2T 1T5
mohit@cs.mcgill.ca

Abstract— For some positive constant ϵ_0 , we give a $(\frac{3}{2} - \epsilon_0)$ -approximation algorithm for the following problem: given a graph $G_0 = (V, E_0)$, find the shortest tour that visits every vertex at least once. This is a special case of the metric traveling salesman problem when the underlying metric is defined by shortest path distances in G_0 . The result improves on the $\frac{3}{2}$ -approximation algorithm due to Christofides [13] for this special case.

Similar to Christofides, our algorithm finds a spanning tree whose cost is upper bounded by the optimum, then it finds the minimum cost Eulerian augmentation (or T-join) of that tree. The main difference is in the selection of the spanning tree. Except in certain cases where the solution of LP is nearly integral, we select the spanning tree randomly by sampling from a maximum entropy distribution defined by the linear programming relaxation.

Despite the simplicity of the algorithm, the analysis builds on a variety of ideas such as properties of strongly Rayleigh measures from probability theory, graph theoretical results on the structure of near minimum cuts, and the integrality of the T-join polytope from polyhedral theory. Also, as a byproduct of our result, we show new properties of the near minimum cuts of any graph, which may be of independent interest.

Keywords-Traveling Salesman Problem; Approximation Algorithms; Randomized Rounding; Random Spanning Trees.

1. INTRODUCTION

The Traveling Salesman Problem (TSP) is a central and perhaps the most well-known problem in combinatorial optimization. TSP has been a source of inspiration and intrigue. In the words of Schrijver [33, Chapter 58], “it belongs to the most seductive problems in combinatorial optimization, thanks to a blend of complexity, applicability, and appeal to imagination”.

In an instance of the TSP, we are given a set of vertices with their pairwise distances and the goal is to find the shortest Hamiltonian cycle which visits every vertex. It is typically assumed that the distance function is a metric.

*Supported by a Stanford Graduate Fellowship. Part of the work was done while the author was a summer intern at IBM Almaden Research Center.

+Supported by National Science Foundation (under grants CCF 0546889, 0729586, and 0915145) and an Alfred Sloan Fellowship.

The best known approximation algorithm for TSP has an approximation factor of $\frac{3}{2}$ and is due to Christofides [13]. Polynomial-time approximation schemes (PTAS) have been found for Euclidean [3], planar [22], [4], [24], or low-genus metrics [15], [14]. On the other hand, the problem is known to be MAX SNP-hard [30] when the distances are one or two (a.k.a (1, 2)-TSP). It is also proved that there is no polynomial-time algorithm with an approximation factor better than $\frac{220}{219}$ for this problem, unless $P = NP$ [29].

The focus of this paper is on traveling salesman problem on *graph metrics*. These are metrics defined by shortest path distances in an arbitrary undirected graph. In other words, we develop an approximation algorithm for the following problem: Given a graph $G_0 = (V, E_0)$, find the shortest tour that visits every vertex at least once.

Theorem 1.1. *The approximation ratio of Algorithm 2 on graph metrics is at most $\frac{3}{2} - \epsilon_0$, where $\epsilon_0 > 0$ is a constant.*

A corollary of our analysis is that the integrality gap of the natural linear programming relaxation (due to Held and Karp [23]) is also strictly below $\frac{3}{2}$ on graph metrics. This ratio is conjectured to be $\frac{4}{3}$ in general. In fact, it has been conjectured that there is a polynomial-time algorithm for TSP with an approximation ratio of $\frac{4}{3}$. This conjecture has been proved in the special case when the underlying graph G_0 is cubic [20], [1], [11].

Recently, and after the first appearance of our paper, Mömke and Svensson [26] came up with a beautiful combinatorial algorithm for the same problem with the approximation ratio of 1.461. Also, An and Shmoys extended the results of this paper to Traveling Salesman Path Problems [2].

1.1. Overview of the Algorithm and Techniques

We propose the same algorithm as in Asadpour et al. [5] for TSP. Let x be the optimum solution of the Held-Karp linear programming relaxation. We sample a tree F from a maximum entropy distribution in which for

every T , $\mathbb{P}[T] \propto \prod_{e \in T} \lambda_e$. We find non-negative λ_e 's in such a way that for every edge $e \in E$ and tree F sampled from μ , $\mathbb{P}[e \in F]$ is proportional to x_e . The details are described in Section 3.

It is not hard to see that the expected cost of the above tree is bounded by the cost of x . We conjecture that the expected cost of the minimum cost Eulerian augmentation of F is strictly less than half of the cost of x for every metric.

Conjecture 1.2. *The approximation ratio of Algorithm 1 on any metric is at most $\frac{3}{2} - \epsilon$, where $\epsilon > 0$ is a constant.*

In this paper, we analyze this algorithm only for graphical metrics and after a slight modification. In our algorithm, we handle the case where x is nearly integral separately using a deterministic algorithm. In fact, when x is nearly integral, it is not hard to find a rounding scheme with an approximation ratio close to $\frac{4}{3}$. When x is not nearly integral, we follow Algorithm 1. See Algorithm 2 for the details.

The analysis of the algorithm has three major ingredients: (i) polyhedral structure of T -join polytope (ii) structure of near minimum cuts, and (iii) properties of random spanning trees. In Part (i), we use the integrality of the T -join polytope to relate the cost of the Eulerian augmentation to the distribution of near minimum cuts and the parity of the edges of F across them. Part (ii) on the structure of near minimum cuts builds on the cactus structure [16] and polygon representation [6] of minimum and near-minimum cuts, respectively. Finally, the last part uses techniques from a recent and very interesting study of strongly Rayleigh measures [9] and their properties to prove results on the joint distribution of the parity of the number of edges across multiple cuts.

Polyhedral structure of T -join polytope. Using the odd-set formulation for the T -join polytope [17], Wolsey [34] showed the cost of T -join for *any* set T of even cardinality is at most *half* of the optimum value of Held and Karp LP. Our starting point is to ask when can the cost of T -join be bounded strictly less than half of Held-Karp bound. To save the cost accounted for an edge e , we identify a sufficient, though not necessary, condition. The condition states that every *near minimum cut* containing edge e , in the graph with weights given by LP solution x , has an even number of edges in the tree selected in the first step.

Structure of Near Minimum Cuts. Let $G(V, E)$ be the weighted (or fractional) graph defined by x . For some δ , consider all $(1 + \delta)$ near minimum cuts or equivalently all cuts of size at most $2(1 + \delta)$. We show that for δ small enough, either a constant fraction of edges appear in a

constant number of $(1 + \delta)$ -near minimum cuts, or x is nearly integral.

In order to build some intuition about the above statement, consider the following extreme case: If x is integral, that is if G is a cycle, then every edge belongs to $n - 1$ minimum cuts where n is the number of vertices. We prove an approximate converse of this statement: for some large constant τ , if almost all the edges are in more than τ near minimum cuts, then the graph is close to a Hamiltonian cycle, in the sense that almost all of its edges are nearly integral.

The above theorem is proved by a characterization of the structure of near minimum cuts for *any* graph and it could be of independent interest. For stating this characterization, we need to define a few things. Let $\mathcal{C}$ be a collection of cuts in graph G . Define a cross graph $\mathcal{G}$ on vertex set $\mathcal{C}$ where an edge between two vertices denotes that their corresponding cuts cross. Every connected component of $\mathcal{G}$ partitions the vertices of G into a set of “atoms”. We show that if $\mathcal{C}$ is a collection of near minimum cuts, the graph resulting from contracting the atoms of any connected component is very close to a cycle. In particular, the weight of nearly all the edges in the resulting graph is very close to half of the size of a minimum cut of G .

Stated in the above form, our result is a generalization of Dinits et al. [16] from minimum cuts to near-minimum cuts. The main technical tool behind the proof is the structure called *polygon representation* of near-minimum cuts as defined by Benczúr [6], [7] and Benczúr and Goemans [8].

Random Spanning Trees and Strongly Rayleigh Measures. In the analysis of this algorithm for asymmetric TSP [5], Asadpour et al. use the negative correlation between the edges of random spanning trees to obtain concentration results on the distribution of edges across a cut. For this work, we have to use stronger virtues of negative dependence [31]. In particular, we use the fact that the distribution of spanning trees belongs to a more general class of measures called Strongly Rayleigh [9]. These measures maintain negative association and log concavity of the rank sequence similar to random spanning trees. In addition, they are closed under projection and truncation and conditioning in certain scenarios.

Let $\mathcal{C}$ be the set of near minimum cuts of G . We prove that for a constant fraction of edges $e \in G$, with constant probability, all of the cuts in $\mathcal{C}$ that contain e have an even intersection with F . Note that the expected number of edges of F across any cut in $\mathcal{C}$ is very close to 2 and it follows simply that a particular cut in $\mathcal{C}$ contains two edges of F with constant probability. Our proof shows the stronger property that with constant

probability, the number of edges of F across all cuts in $\mathcal{C}$ containing e is even. While our proof starts with a λ -uniform spanning tree measure, in intermediate steps we obtain more general measures by conditioning and truncation on certain events. These new measures are no longer λ -uniform spanning tree measures but are still strongly Rayleigh due to closure properties of strongly Rayleigh measures and thus still retain properties like negative association and log concavity. It is instructive to look at the case where this set contains only two degree cuts corresponding to the endpoints of an edge $e = \{u, v\}$. Even in this special case, we are not aware of a direct combinatorial argument to prove that with constant probability, both u and v have an even degree in F .

2. NOTATION AND THE LP RELAXATION

We will use the following linear programming relaxation called LP_{subtour} , known as subtour elimination or Held-Karp linear program. Let $c(\{u, v\})$ denote the distance between u and v or the cost of choosing edge $\{u, v\}$ for each $u, v \in V(G_0)$.

$$\begin{aligned} (LP_{\text{subtour}}) \quad & \text{minimize} \quad \sum_{u, v \in V} c(\{u, v\}) x_{\{u, v\}} \\ \text{subject to} \quad & \sum_{u \in S, v \in \bar{S}} x_{\{u, v\}} \geq 2 \quad \forall S \subsetneq V, \\ & \sum_{u \in V} x_{\{u, v\}} = 2 \quad \forall v \in V, \\ & x_{\{u, v\}} \geq 0 \quad \forall u, v \in V. \end{aligned}$$

With a slight abuse of notation, let x be an optimal solution of this LP. Define $G = (V, E, x)$ to be the fractional support graph corresponding to the optimal vector x , i.e., $E = \{e : x_e > 0\}$.

Throughout the paper, we will refer to x_e as the fraction of edge e in G and to G as a fractional or weighted graph. In this sense, the degree of a node in G is the sum of the fractions of edges incident to that node. Therefore, G is fractionally 2-regular and 2-edge connected.

The following notations will be adopted. For a set $E' \subseteq E$, and any function f defined on the edges of G , let

$$f(E') = \sum_{e \in E'} f(e).$$

For example, $c(E') = \sum_{e \in E'} c(e)$. Similarly, let $x(E') = \sum_{e \in E'} x_e$, and $c(x(E')) = \sum_{e \in E'} c(e)x_e$. In particular, we use $c(x) := c(x(E))$.

For a set $S \subseteq V$, let $E(S) = \{\{u, v\} : u, v \in S\}$ be the set of edges inside S . For two non-crossing sets $S, S' \subset V$, let $E(S, S') = \{\{u, v\} : u \in S, v \in S'\}$ be the set of edges between the vertices in S and S' . In particular, if $S \subset S'$, we use $E(S, S') := \{\{u, v\} : u \in S, v \in S' \setminus S\}$. Also let $\bar{S} = V \setminus S$, and $d(S) = E(S, \bar{S})$ for any $F \subseteq E$.

3. THE ALGORITHM

Our algorithm is quite similar to Christofides algorithm: first it finds a spanning tree whose cost is upper bounded by the optimum, then it finds the minimum cost Eulerian augmentation of that tree.

The main difference is in the selection of the spanning tree. Here, our idea is similar to Asadpour et al. [5]. The algorithm selects a spanning tree randomly from G , the support graph of the solution of LP_{subtour} . The tree is sampled from a distribution μ defined over $\mathcal{T}$, the set of spanning trees of G . This distribution is called λ -uniform or maximum entropy because for every $T \in \mathcal{T}$,

$$\mathbb{P}[T] \propto \prod_{e \in T} \lambda_e.$$

The algorithm finds non-negative λ_e 's in a such a way that for every edge $e \in E$ and tree F sampled from μ , $\mathbb{P}[e \in F]$ is (approximately) equal to $(1 - \frac{1}{n})x_e$. We refer the reader to [5] for more details.

After selecting the spanning tree, the algorithm finds the minimum cost Eulerian augmentation or T -join on the odd-degree vertices of F and constructs a Hamiltonian cycle by short cutting. The details are described in Algorithm 1.

Algorithm 1 Algorithm for TSP for general metrics

Input: A set V of vertices and a cost function $c : V \times V \rightarrow \mathbb{R}^+$ satisfying the triangle inequality.

Output: A hamiltonian tour on V .

- 1: Solve the LP_{subtour} to get an optimum solution x . Let $G = (V, E, x)$ be the support graph of x .
 - 2: Define $z := (1 - 1/n)x$. Let μ denote the maximum entropy distribution over spanning trees of G such that for a spanning tree F sampled from μ , $\mathbb{P}[e \in F] = z_e$ for each edge $e \in E$.
 - 3: Sample a spanning tree F from μ .
 - 4: Let T denote the set of odd-degree nodes in F . Compute the cheapest T -join J .
 - 5: **return** the tour $J \cup F$.
-

In Conjecture 1.2, we conjecture that the expected cost of the tour returned by Algorithm 1 is strictly less than $\frac{3}{2}$ of the cost of OPT for general metrics. However, we can analyze this algorithm only for graphical metrics and after a slight modification. In a special case, where a large fraction of edges in x are nearly integral, we choose the tree deterministically. More specifically, we say an edge $e \in E$ is *nearly integral* if $x_e \geq 1 - \gamma$, where $\gamma > 0$ is a constant. Also x is a nearly integral solution of LP_{subtour} if it has many nearly integral edges, i.e., $|\{e : x_e \geq 1 - \gamma\}| \geq (1 - \epsilon_2)n$ for certain constants $\gamma, \epsilon_2 > 0$. If x is a nearly integral solution of LP_{subtour} , we find the minimum cost spanning tree that contains as many

nearly integral edges as possible. In other words, we find F' the minimum cost spanning subgraph of G_0 that contains all of the nearly integral edges and define F to be the minimum cost spanning tree of F' . Then we simply add minimum T -join on odd-degree vertices of F . The details of our final algorithm are described in Algorithm 2.

Algorithm 2 Improved approximation algorithm for graphic TSP

Input: A set V of vertices and a cost function $c : V \times V \rightarrow \mathbb{R}^+$ satisfying the triangle inequality.

Output: A hamiltonian tour on V .

- 1: Let $\epsilon_2 = 2 \cdot 10^6 \sqrt{\delta}, \gamma = 4 \sqrt{\delta}, \delta = 6.25 \cdot 10^{-16}$.
 - 2: Solve the LP_{subtour} to get an optimum solution x . Let $G = (V, E, x)$ be the support graph of x .
 - 3: **if** x contains $(1 - \epsilon_2)n$ edges of fraction greater than $1 - \gamma$ **then**
 - 4: Find a minimum cost spanning subgraph F' in G_0 that contains all the edges of fraction greater than $1 - \gamma$, and let F be the minimum cost spanning tree in F' .
 - 5: Let T denote the set of odd-degree nodes in F . Compute the cheapest T -join J .
 - 6: **return** the tour $J \cup F$.
 - 7: **else**
 - 8: **return** output of Algorithm 1.
 - 9: **end if**
-

3.1. Overview of the Analysis

In the analysis, we handle the cases considered in Algorithm 1 and Algorithm 2 differently. If x is nearly integral, then a simple polyhedral argument bounds the cost of the tree F and the T -join J (see Lemma 3.3). Indeed the approximation factor is close to $\frac{4}{3}$ in this case.

The more interesting case is when x is not nearly integral, and F is sampled from the distribution μ in Step 3 of Algorithm 1. In that case, first observe that the expected cost of F is at most $c(x)$ since the probability of choosing each edge e is at most x_e . The main part of the argument is to show that the expected cost of the T -join J is smaller than $(1 - \epsilon_0) \frac{c(x)}{2}$.

$$\begin{aligned}
 (LP_{T\text{-join}}) \quad & \text{minimize} \quad \sum_{e \in E} c(e) y_e \\
 \text{subject to} \quad & \sum_{e \in E(S, \bar{S})} y_e \geq 1 \quad \forall S \subseteq V, |S \cap T| \text{ odd}, \\
 & y_e \geq 0 \quad \forall e \in E.
 \end{aligned}$$

In order to bound the cost of the T -join, first observe that half of any solution of LP_{subtour} , the vector $\frac{x}{2}$, is a feasible fractional solution to the $LP_{T\text{-join}}$ for any set $T \subseteq V$. This is because across any cut, the sum of

the fractions of $\frac{x}{2}$ is at least 1. This observation, made originally by Wolsey [34], also implies that the solution of Christofides is at most $\frac{3}{2}c(x)$.

In order to get a factor better than $\frac{3}{2}$, it is sufficient to construct a feasible solution of smaller cost for the T -join polytope, when T is the set of odd degree vertices of the sampled spanning tree F . When T in $LP_{T\text{-join}}$ is set to the odd-degree vertices of F , the constraints present are exactly for the cuts which intersect in odd number of edges with F .

A cut is a $(1 + \delta)$ near minimum cut of G if the total fraction of the edges in the cut is at most $(1 + \delta)$ times the minimum cut of G . In other words, cuts $(S, \bar{S})$ for which $x(E(S, \bar{S})) \leq 2(1 + \delta)$ are called near minimum cuts. Also, a cut $(S, \bar{S})$ is odd with respect to F iff $F \cap E(S, \bar{S})$ is odd, i.e., F contains an odd number of edges of the cut $(S, \bar{S})$. The following two definitions are crucial.

We say an edge e is **even** with respect to F if any near minimum cut that includes e is even with respect to F , i.e., for all $(S, \bar{S})$ such that $e \in E(S, \bar{S})$ and $x(E(S, \bar{S})) \leq 2(1 + \delta)$, $|F \cap E(S, \bar{S})|$ is even. Given a tree F , setting $y_e = \frac{x_e}{2(1 + \delta)}$ for each edge e which is even with respect to F and $y_e = \frac{x_e}{2}$ for every other edge e , we obtain a feasible solution to the $LP_{T\text{-join}}$ when T is the set of odd-degree vertices of F . Thus it is enough to find a tree F for which the set of even edges is large.

Let $\mathcal{E}(e)$ be the event that e is even with respect to F where F is sampled from the distribution μ . We say e is **good** if the probability of this event is bounded from zero by some constant. More precisely, if for a fixed constant $\rho > 0$,

$$\begin{aligned}
 \mathbb{P}[\exists (S, \bar{S}) : e \in E(S, \bar{S}) \text{ and } x(E(S, \bar{S})) \leq 2(1 + \delta) \\
 \text{and } |F \cap E(S, \bar{S})| \text{ is odd}] \leq 1 - \rho.
 \end{aligned}$$

Our strategy is to identify a large number of good edges in the graph. We will use these edges to show that the cost of T -join is strictly less than $\frac{c(x)}{2}$. The following Theorem shows that it is indeed possible to find such edges if the algorithm samples the tree F in Step 3.

Theorem 3.1 (Structure Theorem). *Let x be an optimal solution of LP_{subtour} , and let μ be the λ -uniform measure defined based on x . There exist sufficiently small constants ϵ_1, ρ bounded away from zero such that at least one of the following is satisfied by x :*

- 1) **there is an abundance of good edges in x :** There exists a set $E^* \subseteq E$ such that $x(E^*) \geq \epsilon_1 n$, and
$$\forall e \in E^* : \mathbb{P}[\mathcal{E}(e)] \geq \rho.$$
- 2) **x is nearly integral:** x contains at least $(1 - \epsilon_2)n$ edges of fraction greater than $1 - \gamma$.

We note that the Structure Theorem is valid for *all* feasible solutions to the Held-Karp relaxation and not

just for vertex solutions of the linear program which have been studied extensively [12], [10], [21].

3.2. Proof of Theorem 1.1

In the rest of the section, we show Theorem 3.1 implies Theorem 1.1 by constructing feasible solutions to $LP_{T\text{-join}}$ of small cost.

Case 1: x has at least $\epsilon_1 n$ good edges. First observe that the expected cost of F is at most $c(x)$. This is because $\mathbb{P}[e \in F] = z_e \leq x_e$ so we have $\mathbb{E}[c(F)] = \sum_{e \in E^*} c(e) \mathbb{P}[e \in F] \leq c(x)$. Hence, we only need to bound the cost of the T -join.

Lemma 3.2. *Let x be a fractional solution of (LP_{subtour}) , $E^* \subset E$ be the set of good edges. If there are a lot of good edges, that is if $x(E^*) \geq \epsilon_1 n$, then the expected cost of the smallest Eulerian tour is at most $3/2 - \frac{\epsilon_1 \delta \rho}{4(1+\delta)}$.*

Proof. We provide a fractional solution to the $(LP_{T\text{-join}})$ to make it Eulerian. For any edge $e \in E$ if e is contained in at least one odd $(1+\delta)$ near minimum cut $(S, \bar{S})$, set $y_e = x_e/2$, otherwise set $y_e = x_e/2(1+\delta)$. Observe that a cut $(S, \bar{S})$ is odd in F iff $|S \cap T|$ is odd. Therefore, y is indeed a fractional solution of $(LP_{T\text{-join}})$. Now, to bound the cost of y in Step 3 note:

$$\begin{aligned} \mathbb{E}[c(y)] &\leq \frac{c(x)}{2} - \sum_{e \in E} x_e c(e) \mathbb{P}[e \text{ is even}] \left(\frac{1}{2} - \frac{1}{2(1+\delta)} \right) \\ &\leq \frac{c(x)}{2} - \frac{\delta}{2(1+\delta)} \sum_{e \in E^*} x_e \rho \leq c(x) \left(\frac{1}{2} - \frac{\epsilon_1 \delta \rho}{4(1+\delta)} \right). \end{aligned}$$

The second inequality holds because $c(e) \geq 1$ for all $e \in E$, and the last one because $c(x) \leq 2n$. Since the T -join polytope is integral [17], the minimum cost integral T -join costs at most $c(y)$. By adding the edges of minimum T -join J to F we obtain an Eulerian tour of expected total weight at most $c(x) \left(\frac{3}{2} - \frac{\epsilon_1 \delta \rho}{4(1+\delta)} \right)$. $\square$

The above argument bounds the cost of the tour in expectation. By sampling a tree $\log n$ times and choosing the best solution, one can obtain an Eulerian tour of the desired cost with high probability.

Case 2: x is nearly integral. In this case, we bound the cost of the tree F and T -join J together and prove the following lemma. The construction of the fractional T -join in the lemma is similar to a construction by Monma, Munson and Pulleyblank [27].

Lemma 3.3. *Let x be a fractional solution of (LP_{subtour}) . If x contains at least $(1 - \epsilon_2)n$ edges of fraction greater than $1 - \gamma$, then the tour computed in Algorithm 2, step 5 is at most $c(x) \left(\frac{4}{3} + 2\epsilon_2 + 4\gamma \right)$.*

Proof. Let $I' = \{e \mid x_e > 1 - \gamma\}$ be the set of nearly integral edges, and let F' be the minimum cost spanning

graph that contains I' . Since G_0 is connected, I' can be augmented into a connected graph using only edges of cost 1. Hence, we have $c(F') = c(I') + |F' \setminus I'| \leq \frac{\sum_{e \in I'} c(e) x_e}{1-\gamma} + |F' \setminus I'|$.

Recall that F is a minimum cost spanning subgraph of F' . Because of the constraints of LP and since $\gamma < 1/3$, it is easy to see that I' consists of disjoint cycles and paths and the length of each cycle is at least $\frac{1}{\gamma}$. Therefore, F will have at least $n(1 - \epsilon_2)(1 - \gamma)$ edges from I' . Therefore, $|F \setminus I'| \leq n(\epsilon_2 + \gamma)$. Let us set $I = I' \cap F$.

Let T denote the set of odd vertices in F . Again, we bound the cost of T -join by constructing a fractional solution to the $LP_{T\text{-join}}$, and then invoking the integrality of the T -join polytope.

Let $y_e = \frac{x_e}{3(1-\gamma)}$ for $e \in I$, $y_e = 1$ for $e \in F \setminus I$, and $y_e = x_e$ otherwise. We first show that y is feasible for $LP_{T\text{-join}}$. Let $(U, \bar{U})$ be any cut which has an odd number of vertices of T in U (equivalently, a cut that has an odd number of edges of F). If there exists an $e \in (F \setminus S) \cap E(U, \bar{U})$, then $y(d(U)) \geq y_e \geq 1$ and the constraint is satisfied. Otherwise, we have $E(U, \bar{U}) \cap F \subseteq S$. Therefore since $(U, \bar{U})$ has an odd number of edges F , and $I \subset F$, $(U, \bar{U})$ must contain an odd number of edges of I . By the values assigned to the edges in y , we have

$$\begin{cases} y(d(U)) \geq x(d(U) \setminus I) \geq 1 & \text{if } |I \cap E(U, \bar{U})| = 1 \\ y(d(U)) \geq y(d(U) \cap I) \geq 1 & \text{if } |I \cap E(U, \bar{U})| \geq 3 \end{cases}$$

thus y is a feasible solution of $(LP_{T\text{-join}})$.

Now we bound the cost of the final Eulerian subgraph which will be at most $c(F) + c(y)$

$$\begin{aligned} c(F) + c(y) &\leq \frac{c(x(I))}{1-\gamma} + 2c(F \setminus I) + c(x(E \setminus F)) + \frac{c(x(I))}{3(1-\gamma)} \\ &\leq \frac{4c(x(I))}{3(1-\gamma)} + n(2\epsilon_2 + 2\gamma) + c(x(E \setminus I)) \\ &\leq c(x) \left(\frac{4}{3} + 2\epsilon_2 + 4\gamma \right). \end{aligned}$$

The last inequality follows from the fact that $c(x) \geq n$, and the last inequality follows from $\gamma < 1/3$. $\square$

4. PROOF OF THE STRUCTURE THEOREM (MAIN IDEAS)

The rest of the paper is dedicated to proving the Structure Theorem. For proving this theorem, we have to establish several results about the structure of near minimum cuts in graphs as well as properties of random spanning trees. Because of the space constraint we do not include most of the proofs. We will explain the main ideas instead.

Subsections 4.1 and 4.2 focus on the structure of near minimum cuts. The following two definitions are quite important.

Definition 4.1 (Atom). For a collection $\mathcal{C}$ of cuts of a graph $G = (V, E)$, atoms of $\mathcal{C}$ are the members of a partition $\mathcal{P}$ of the vertex set V such that

- no cut of $\mathcal{C}$ divides any of the atoms of $\mathcal{C}$, and
- $\mathcal{P}$ is the coarsest partition with this property.

We say an atom is **singleton** if it is a set of a single vertex of V .

Definition 4.2 (Cross Graph). A pair of cuts $(A, \bar{A})$ and $(B, \bar{B})$ is said to cross if $A \cap B, A \setminus B, B \setminus A, V \setminus (A \cup B)$ are all non-empty. For a collection $\mathcal{C}$ of cuts of a graph $G = (V, E)$, **cross graph** $\mathcal{G}$ is a graph on vertex set $\mathcal{C}$ and that has an edge between two cuts in $\mathcal{C}$ if they cross. Each connected component of $\mathcal{G}$ is called a **cut class**.

Subsection 4.1, reduces the proof of Theorem 3.1 to the following: either x is nearly integral, or a constant fraction of edges of x are contained in a constant number of near minimum cuts. More precisely, if the number of atoms in large cut classes defined by $(1 + \delta)$ near minimum cuts of G is close to n , then case 2 of Theorem 3.1 holds. Otherwise, a constant fraction of edges of G are not contained in any large cut class. This is stated in Lemmas 4.4 and 4.6. The proof of these lemmas build on a characterization of near minimum cuts of G presented in subsection 4.2. We note that this characterization, presented in Theorem 4.9, applies to *any* graph and could be of independent interest.

Finally, in subsection 4.3, we identify E^* from the edges that are contained only in constant number of near minimum cuts.

4.1. Atoms, Cut Classes and Cross Graphs

Consider the cross graph corresponding to $(1 + \delta)$ -near minimum cuts of G and let $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_l$ be its cut classes. Denote the set of atoms of any of these families of cuts by $\phi(\mathcal{C}_i)$ for $1 \leq i \leq l$.

Definition 4.3. Let $\tau = \frac{1}{20\sqrt{\delta}} = 2 \cdot 10^6$. We say a cut class $\mathcal{C}_i$ is **large** if $|\phi(\mathcal{C}_i)| \geq \tau$, and **small** otherwise.

Let $L(\tau)$ be the set of all atoms of the large cut classes, i.e.

$$L(\tau) = \bigcup_{\mathcal{C}_i: |\phi(\mathcal{C}_i)| \geq \tau} \phi(\mathcal{C}_i).$$

The size of $L(\tau)$ plays an important role. It is easy to see that $|L(\tau)| \leq n(1 + \frac{2}{\tau-2})$. Now, if $|L(\tau)|$ is close to its maximum possible value, i.e. $|L(\tau)| \geq (1 - \epsilon)n$, then case 2 of Theorem 3.1 holds.

Lemma 4.4. For any $\epsilon \geq \frac{1}{\tau-2}$, and $\delta < \frac{1}{100}$, if $|L(\tau)| \geq (1 - \epsilon)n$ then G contains at least $(1 - 20\sqrt{\delta} - 17\epsilon)n$ edges of fraction greater than $1 - 4\sqrt{\delta}$.

In order to understand this intuitively, think about the cross graph defined by the minimum cuts of a cycle

of length n . Observe that this graph contains $\binom{n}{2}$ (near) minimum cuts, and the cross graph $\mathcal{G}$ contains only one connected component or equivalently, one large cut class with n atoms. Therefore, if G is a cycle of length n , $|L(\tau)| = n$. The above Lemma is an approximate converse of this observation: if $|L(\tau)|$ is large, then the LP solution is in a certain way close to a Hamiltonian cycle.

On the other hand, if $L(\tau)$ is small, the first case of Theorem 3.1 holds. We show this in two steps. The first step shows that there exists a large fraction of edges contained in a constant number of near minimum cuts. A necessary condition for this property is that the edge does not belong to any large cut classes.

Definition 4.5. An edge e is incident to an atom a , if exactly one of its endpoints is contained in a . An edge e is said to be contained in a cut class $\mathcal{C}_i$ if e is incident to some atom of $\mathcal{C}_i$.

Let E_S be the set of edges that are not contained in any of the large cut classes. In the next lemma we show that if $|L(\tau)| < (1 - \epsilon)n$, then $x(E_S)$ is large:

Lemma 4.6. If $|L(\tau)| < (1 - \epsilon)n$ then $x(E_S) \geq n(\epsilon - 3\delta)$.

A simple double counting argument shows that a constant fraction of edges in E_S are contained in only a small number of near minimum cuts. However, appearing in a constant number of near minimum cuts does not automatically guarantee that an edge is good (see Figure 1 for a counter example). We will find the set of good edges E^* in E_S in Section 4.3.

Before that, we are ready to assign the exact values of the constants. We set $\epsilon = \frac{5000}{\tau} = 2.5 \cdot 10^{-3}$ so as to satisfy all the conditions. This already implies appropriate values for ϵ_1, ϵ_2 and γ in the algorithm. We set $\epsilon_1 = 3000\delta$, $\epsilon_2 = 2 \cdot 10^6 \sqrt{\delta} \geq 20\sqrt{\delta} + 17\epsilon$ and $\gamma = 4\sqrt{\delta}$. Finally, from Lemma 3.3, $2\epsilon_2 + 4\gamma \leq 0.11$ is enough to give a better than $\frac{3}{2}$ bound on the performance of the algorithm. This implies $\delta = 6.25 \cdot 10^{-16}$ suffices to satisfy all the conditions. We note that we have not optimized the constants.

4.2. Near Minimum Cuts and their Cactus-like Structure

In this section, we prove crucial lemmas about the structure of near minimum cuts of any graph. Applying these lemmas to the solution of the Held-Karp linear program directly yields Lemma 4.4 and Lemma 4.6.

Let H be an unweighted graph and let c denote size of the minimum cut of H . For a partitioning $\mathcal{P} = \{P_1, P_2, \dots, P_k\}$ of vertices in H , let $H(\mathcal{P})$ be the graph obtained by identifying the vertex set of each part P_i , and removing the self-loops afterwards. For example, for a cut class $\mathcal{C}_i$, each vertex of $H(\phi(\mathcal{C}_i))$ is corresponding to an atom of $\mathcal{C}_i$.

The following lemma about the structure of minimum cuts follows from the cactus representation [16] (also see Fleiner and Frank [19] for a short proof).

Lemma 4.7. [16] *Let C_i denote a cut class of minimum cuts of H . Then $H(\phi(C_i))$ is a cycle where weight of every edge is exactly $\frac{\epsilon}{2}$ and every pair of edges of the cycle corresponds to a minimum cut of H .*

Our main result in this section is that the above lemma generalizes to near minimum cuts in an approximate sense.

Definition 4.8 $((\alpha, \alpha', \beta)$ -cactaceous). *A graph $H = (V, E)$ with minimum cut c is (α, α', β) -cactaceous if for some $\delta \geq 0$:*

- *There exists at least $m := (1 - \alpha\sqrt{\delta})|V(H)|$ pairs of vertices of H , $\{(v_1, u_1), (v_2, u_2), \dots, (v_m, u_m)\}$ such that for each $1 \leq i \leq m$, $E(v_i, u_i) \geq \frac{\epsilon}{2}(1 - \alpha'\sqrt{\delta})$, and each vertex $v \in V(H)$ is contained in at most two such pairs.*
- *The number of edges of the graph H satisfies the following:*

$$\frac{\epsilon}{2}|V(H)| \leq |E(H)| \leq (1 + \beta\delta)\frac{\epsilon}{2}|V(H)|.$$

Theorem 4.9. *For any $\delta < 1/100$, let C_i denote a cut class of $(1 + \delta)$ near minimum cuts of H . Then $H(\phi(C_i))$ is $(20, 4, 3)$ -cactaceous.*

If we let $\delta = 0$ in the description of Theorem 4.9, we obtain that for any cut class C_i of the collection of minimum cuts of H , the graph $H(\phi(C_i))$ is a cycle where the weight of each edge is $\frac{\epsilon}{2}$, thus we obtain Lemma 4.7.

The main technical tool behind the proof is the structure called *polygon representation* of near-minimum cuts as defined by Benczúr [6], [7], and Benczúr, Goemans [8]. Benczúr showed that for $\delta \leq 1/5$, the near minimum cuts of any graph H can be represented using polygon representation. Our theorem uses this representation heavily. However, the emphasis of [8] (and results before that) were on representing the vertex sets of minimum cuts. Instead, here we focus on the edge sets and observe several interesting properties that could be of independent interest.

Applying the above results to the structure of near minimum cuts of the Held-Karp linear programming solution x leads to the proofs of Lemmas 4.6 and 4.4. Firstly, Theorem 4.9 implies the following corollary about the structure of any cut class of near minimum cuts of the weighted graph $G = (V, E, x)$.

Corollary 4.10. *For any $\delta < 1/100$, let C_i be a cut class of the $(1 + \delta)$ near minimum cuts of the weighted graph $G = (V, E, x)$. Then $G(\phi(C_i))$ satisfies the following:*

- *There exists at least $m := (1 - 20\sqrt{\delta})|\phi(C_i)|$ pairs of vertices of $G(\phi(C_i))$, $\{(a_1, b_1), (a_2, b_2), \dots, (a_m, b_m)\}$*

such that for each $1 \leq i \leq m$, $x(a_i, b_i) \geq 1 - 4\sqrt{\delta}$, and each vertex $a \in V(G(\phi(C_i)))$ is contained in at most two such pairs.

- $|\phi(C_i)| \leq x(E(G(\phi(C_i)))) \leq (1 + 3\delta)|\phi(C_i)|$.

Proof of Lemma 4.6. We will show that a constant fraction of edges in G are not incident to any of the atoms of $L(\tau)$. By Corollary 4.10, for any cut class C_i , the total sum of the fraction of edges in $G(\phi(C_i))$ is at most $|\phi(C_i)|(1 + \beta\delta)$, where $\beta := 3$. Thus the total fraction of edges that are contained in at least one of the large cut classes is no more than

$$|L(\tau)|(1 + \beta\delta) < n(1 - \epsilon)(1 + \beta\delta) \leq n - n(\epsilon - \beta\delta).$$

Therefore, since the total sum of the fraction of edges in G is n (i.e. $x \in LP_{\text{subtour}}$), we have $x(E_S) \geq (\epsilon - \beta\delta)n$. $\square$

Proof of Lemma 4.4. We will show that for any $\delta < \frac{1}{100}$, if $|L(\tau)| \geq (1 - \epsilon)n$, then x contains at least $n(1 - \alpha\sqrt{\delta} - 5\epsilon - \frac{12}{\tau-2}) \geq n(1 - \alpha\sqrt{\delta} - 17\epsilon)$ edges of fraction $1 - \alpha'\sqrt{\delta}$, where $\alpha := 20, \alpha' := 4$. This is because if $|L(\tau)|$ is large, then most of the atoms in $L(\tau)$ are singletons. Also, a near integral edge in $G(\phi(C_i))$ incident to two singleton atoms corresponds to an actual near integral edge in G .

Let L be the number of large cut classes. Using the “tree hierarchy” of minimum cuts defined by Benczúr [7] one can simply prove the following claim (the claim is proved in the full version of the paper [28]):

Claim 4.11. *Any subset $C_1, C_2, \dots, C_l$ of the cut classes of the cross graph $\mathcal{G}$, contains a set of $\sum_{i=1}^l |\phi(C_i)| - 2(l-1)$ mutually disjoint atoms.*

Since the number of mutually disjoint subsets of a set of n vertices is no more than n , we have $|L(\tau)| \leq n + 2L$. But then we have $L\tau \leq n + 2L$ and therefore $L \leq \frac{n}{\tau-2}$. Therefore, we can find at least $|L(\tau)| - 2L \geq n(1 - \epsilon - \frac{2}{\tau-2})$ mutually disjoint atoms in $L(\tau)$. But these atoms define a partition of the ground set V , and at least $n(1 - 2\epsilon - \frac{4}{\tau-2})$ of them must be singletons. Therefore, the number of non-singleton atoms of $L(\tau)$ is at most $n(2\epsilon + \frac{6}{\tau-2})$.

On the other hand, by Corollary 4.10, there are

$$\sum_{C_i: |\phi(C_i)| \geq \tau} |\phi(C_i)|(1 - \alpha\sqrt{\delta}) = |L(\tau)|(1 - \alpha\sqrt{\delta}) \geq n(1 - \epsilon - \alpha\sqrt{\delta})$$

edges of fraction $1 - \alpha'\sqrt{\delta}$ in graphs $G(\phi(C_i))$ for any large cut class C_i . Hence, at least $n(1 - \alpha\sqrt{\delta} - 5\epsilon - \frac{12}{\tau-2})$ of these edges are incident only to singletons. But edges adjacent to two singletons are corresponding to actual edges of G . We conclude that there are $n(1 - \alpha\sqrt{\delta} - 5\epsilon - \frac{12}{\tau-2})$ edges of fraction $1 - \alpha'\sqrt{\delta}$ in G . Then lemma follows from the assumption $\epsilon > \frac{1}{\tau-2}$. $\square$




Figure 1. The left diagram represents the λ values of the edges, while the right diagram represents the approximate probability of each edge. The example shows that although $\mathbb{P}[(u, v) \in T] \approx \frac{1}{2}$, and the expected degree of u and v is 2, $\mathbb{P}[\deg_T(u) + \deg_T(v) = 3] = 1 - o(1)$. Therefore (u, v) is not good. Note that this is not exactly a solution to LP_{subtour} , but still points to difficulties in proving a trivial edge is good.

4.3. Good Edges and Random Spanning Trees

In this section, we prove that at least a constant fraction of the edges in E_S (as defined in Lemma 4.6) are good.

Theorem 4.12. *If $|L(\tau)| \leq (1 - \epsilon)n$ (i.e., x is not nearly integral), then there exists a set $E^* \subseteq E_S$ of good edges such that $x(E^*) \geq \epsilon_1 n$.*

We identify three types of good edges, namely, “trivial edges”, “inside edges”, and “thread edges”. Here, we only define trivial edges and inside edges. Then we prove that certain trivial edges are good. We refer the reader to the full version of this paper [28] for the complete proof.

Definition 4.13 (Trivial Edge). *We call an edge trivial if it is contained in only two near minimum cuts, which are the degree constraint of its endpoints.*

Definition 4.14 (Inside Edge). *An edge $e = (u, v)$ is an inside edge if it is contained in a small cut class C_i such that the atoms of C_i containing u and v are singletons.*

Our main tools for proving that an edge $e \in E_S$ is good are properties of strongly Rayleigh measures. Strongly Rayleigh measures satisfy “the strongest form” of negative dependence [31] and are studied in great details in Borcea et al. [9]. These measures include uniform and λ -uniform random spanning tree measures as a special case.

Let us define these measures and describe some of the main properties that will be used in our proof. For an element (edge) $e \in E$, let X_e be the indicator random variable of e , and for $S \subseteq E$, let $X_S = \sum_{e \in S} X_e$. For any measure μ on 2^E , we may form its generating polynomial, $f(t) = \sum_{S \subseteq E} \mu(S) t^S$, where $t^S = \prod_{e \in S} t_e$. A polynomial $f \in \mathbb{C}[t_e]_{e \in E}$ is called *stable* if $f(t_e : e \in E) \neq 0$ whenever $\text{Im}(t_e) > 0$ for all $e \in E$. A stable polynomial with

all real coefficients is called *real stable*. For example, a polynomial in one variable is real stable iff all its roots are real. A measure μ on 2^E is called *strongly Rayleigh* if its generating polynomial is real stable.

Strongly Rayleigh measures are closed under Projection, Conditioning, and Truncation, and they satisfy Negative Association, Ultra Log Concavity of the rank sequence, and Stochastic Domination Property on truncated measures.

- **Truncation:** Let μ be a probability measure on 2^E , and $1 \leq k \leq |E|$. The *truncation* of μ to k is the conditional measure $\mu_k := (\mu \mid X_E = k)$. Borcea et al. [9] showed that Strongly Rayleigh measures are closed under truncation.
- **Negative Association:** A measure μ on 2^E is called *negatively associated* or NA if

$$\mathbb{E}_\mu[f] \mathbb{E}_\mu[g] = \int f d\mu \int g d\mu \geq \int f g d\mu = \mathbb{E}_\mu[fg]$$

for any increasing functions f, g on 2^E that depend on **disjoint** sets of elements.

Feder and Mihail [18] proved that uniform measures on balanced matroids (and in particular on spanning trees) satisfy negative association. Borcea et al. [9] proved that the strongly Rayleigh measure also satisfy the negative association property.

- **Log Concavity:** A real sequence $\{a_k\}_{k=0}^m$ is LC if $a_k^2 \geq a_{k-1}a_{k+1}$, $1 \leq k \leq m-1$, and the indices of its non-zero terms form an interval (of non-negative integers).

Borcea et al. [9] showed that the rank sequence of a strongly Rayleigh measure is the probability distribution of the number of successes in m independent trials for some sequence of success probabilities

$p_1, \dots, p_m$. Therefore, it satisfies the LC property [32].

- **Stochastic Domination:** Let μ, ν be two measures defined on 2^E . We say μ *stochastically dominates* ν ($\nu \leq \mu$) if for any increasing event $\mathcal{A}$ on 2^E , we have $\mu(\mathcal{A}) \geq \nu(\mathcal{A})$.

Borcea et al. [9] showed that a truncation of strongly Rayleigh measures is stochastically dominated by a truncation of a larger value, i.e., for all $1 \leq k < m$, $\mu_k \leq \mu_{k+1}$.

In the rest of the paper, we use the properties of strongly Rayleigh measures to prove that certain types of trivial edges are good. Trivial edges are the simplest possible candidate edges for being good. In fact, in the extended version of this paper [28], we show that any trivial edge e such that $x_e < \frac{1}{2} - \frac{1}{8000}$ or $x_e > \frac{1}{2} + \frac{1}{8000}$ is good. Furthermore, of any adjacent pair of trivial edges, one of them is good. The reader can see Figure 1 for an example of a trivial edge which is not good. In the following lemma, we prove this claim when x_e is very small to illustrate the techniques used.

Lemma 4.15. *Let μ be a λ -uniform measure of spanning trees of $G = (V, E, x)$. There exists a small but fixed $\epsilon > 0$ such that for any trivial edge $e = (u, v)$, if $x_e \leq \epsilon$ then e is even with a constant probability (i.e., u and v have even degree with a constant probability).*

Proof. Since $x_e \leq \epsilon$ and ϵ is very small, it is sufficient to condition on the spanning trees that does not contain e , and show that u and v have an even degree with a constant probability in the conditional distribution. We let $\nu = \{\mu : X_e = 0\}$ where X_e is the indicator random variable of edge e . Observe that ν is a strongly Rayleigh measure since μ is strongly Rayleigh. Actually, it is not difficult to see that ν is also a uniform spanning tree measure. By the negative association property of the uniform spanning tree measure, this conditioning does not change the expected degree of e 's endpoints by more than ϵ . Therefore, we have that $2 - \epsilon \leq \mathbb{E}_\nu[|T \cap \deg(u)|], \mathbb{E}_\nu[|T \cap \deg(v)|] \leq 2$.

Let $Y := |T \cap d(u)|$, $Z := |T \cap d(v)|$. Since edges are negatively correlated under λ -uniform measures of spanning trees [25], the degree of both u and v are highly concentrated around their expected value. It follows that $\mathbb{P}_\nu[Y = 2]$ and $\mathbb{P}_\nu[Z = 2]$ are constants. Since these constants may be strictly below $1/2$, we may not simply apply a union bound to prove the lemma. Instead, we show that $\mathbb{P}_\nu[Z = 2 \mid Y = 2]$ is a constant. Note that the measure $\{\nu \mid Y = 2\}$ is not necessarily a λ -uniform spanning tree measure, but it is still a strongly Rayleigh measure since it is a truncation of a strongly Rayleigh measure. We use concentration properties of this measure to prove the lemma.

Let $E' := E \setminus d(u)$, and ν' be the projection of ν onto E' . Since any spanning tree has exactly $n - 1$ edges in G , the distribution of Z under ν conditioned on $Y = k$ is the same as the distribution of Z under ν'_{n-1-k} (recall that since $X_e = 0$ under ν , $d(v) \subseteq E'$). Therefore, it is sufficient to prove that $\mathbb{P}_{\nu'_{n-3}}[Z = 2]$ is a constant. Moreover, since ν is a strongly Rayleigh measure, and these measures are closed under projection and truncation, ν'_{n-3} is strongly Rayleigh. Hence, ν'_{n-3} satisfies negative association and therefore Z is highly concentrated around its expected value. Thus, it is sufficient to show that $\mathbb{E}_{\nu'_{n-3}}[Z]$ is strictly above 1 and below 3.

First, we show $\mathbb{E}_{\nu'_{n-2}}[Z] \leq 3$, then we use stochastic domination to prove the lemma. Let $\mathcal{A} := \{Y \leq 1\}$. Since $\mathcal{A}$ is a decreasing event, and for any edge $f \in E'$, X_f is an increasing function, by the negative association property $\mathbb{P}[X_f | \mathcal{A}] \geq \mathbb{P}[X_f]$. Since conditioning on $\mathcal{A}$ implies $Y = 1$ and $X_{E'} + Y = n - 1$, we have $\mathbb{E}_\nu[X_{E'} | \mathcal{A}] = \mathbb{E}_\nu[X_{E'}] + 1$; thus

$$\mathbb{E}_{\nu'_{n-2}}[Z] = \mathbb{E}_\nu[Z | \mathcal{A}] \leq \mathbb{E}_\nu[Z] + 1 = 3.$$

Therefore, by the stochastic domination, we obtain

$$1 \leq \mathbb{E}_{\nu'_{\leq n-4}}[Z] \leq \mathbb{E}_{\nu'_{n-3}}[Z] \leq \mathbb{E}_{\nu'_{n-2}}[Z] \leq 3, \quad (1)$$

where $\mathbb{E}_{\nu'_{\leq n-4}}[Z]$ is the measure ν' conditioned on $X_{E'} \leq n - 4$, and the first inequality simply follows by the connectivity property of spanning trees (i.e., we always have $Z \geq 1$).

On the other hand, since $\mathbb{E}_\nu[Z] \simeq 2$ we have

$$\begin{aligned} \mathbb{P}_\nu[Y = 1] \mathbb{E}_{\nu'_{n-2}}[Z] &+ \mathbb{P}_\nu[Y = 2] \mathbb{E}_{\nu'_{n-3}}[Z] \\ &+ \mathbb{P}_\nu[Y \geq 3] \mathbb{E}_{\nu'_{\leq n-4}}[Z] \simeq 2. \end{aligned}$$

But if $\mathbb{E}_{\nu'_{n-3}}[Z]$ is very close to 1 (or 3), then we must have $\mathbb{P}_\nu[Y = 1] \geq \frac{1}{2}$ (resp. $\mathbb{P}_\nu[Y \geq 3] \geq \frac{1}{2}$). This is in contradiction with the expected value of Y being equal to 2. In other words, since $\mathbb{E}_\nu[Y] \simeq 2$, a simple application of Chernoff-Hoeffding bounds implies that $\mathbb{P}_\nu[Y = 1]$ and $\mathbb{P}_\nu[Y \geq 3]$ are strictly less than $\frac{1}{2}$. Therefore, $\mathbb{P}_{\nu'_{n-3}}[Z] = \mathbb{P}_\nu[Z = 2 \mid Y = 2]$ is a constant, and e is even with a constant probability. $\square$

ACKNOWLEDGMENT

We would like to thank Persi Diaconis, Russel Lyons, Elchanan Mossel, Robin Pemantle, Mohammad Reza Salavatipour and Morteza Zadimoghaddam for very helpful discussions and pointers to the literature. We also would like to thank Hyung-Chan An, Mohammad-Hossein Bateni and Moses Charikar for a careful reading of the paper and extensive comments.

REFERENCES

- [1] N. Aggarwal, N. Garg, and S. Gupta, "A 4/3-approximation for tsp on cubic 3-edge-connected graphs," *CoRR*, vol. abs/1101.5586, 2011.

- [2] H.-C. An and D. B. Shmoys, "Lp-based approximation algorithms for traveling salesman path problems," *CoRR*, vol. abs/1105.2391, 2011.
- [3] S. Arora, "Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems," in *JACM*, 1996, pp. 2–11.
- [4] S. Arora, M. Grigni, D. Karger, P. Klein, and A. Woloszyn, "A polynomial-time approximation scheme for weighted planar graph tsp," in *SODA*, 1998, pp. 33–41.
- [5] A. Asadpour, M. X. Goemans, A. Madry, S. Oveis Gharan, and A. Saberi, "O(logn/loglogn) approximation to the asymmetric traveling salesman problem," *SODA*, 2010.
- [6] A. A. Benczúr, "A representation of cuts within 6/5 times the edge connectivity with applications," in *FOCS*, 1995, pp. 92–102.
- [7] A. A. Benczur, "Cut structures and randomized algorithms in edge-connectivity problems," Ph.D. dissertation, 1997.
- [8] A. A. Benczúr and M. X. Goemans, "Deformable polygon representation and near-mincuts," *Building Bridges: Between Mathematics and Computer Science, M. Groetschel and G.O.H. Katona, Eds., Bolyai Society Mathematical Studies*, vol. 19, pp. 103–135, 2008.
- [9] J. Borcea, P. Branden, and T. M. Liggett, "Negative dependence and the geometry of polynomials," *Journal of American Mathematical Society*, vol. 22, pp. 521–567, 2009.
- [10] S. C. Boyd and W. R. Pulleyblank, "Optimizing over the subtour polytope of the travelling salesman problem," *Math. Program.*, vol. 49, pp. 163–187, 1991.
- [11] S. Boyd, R. Sitters, S. van der Ster, and L. Stougie, "Tsp on cubic and subcubic graphs," in *IPCO*, 2011.
- [12] R. D. Carr and S. Vempala, "On the held-karp relaxation for the asymmetric and symmetric traveling salesman problems," *Math. Program.*, vol. 100, no. 3, pp. 569–587, 2004.
- [13] N. Christofides, "Worst case analysis of a new heuristic for the traveling salesman problem," Graduate School of Industrial Administration, Carnegie-Mellon University, Pittsburgh, PA, Report 388, 1976.
- [14] E. D. Demaine, M. Hajiaghayi, and K. ichi Kawarabayashi, "Contraction decomposition in h-minor-free graphs and algorithmic applications," in *STOC*, 2011.
- [15] E. D. Demaine, M. Hajiaghayi, and B. Mohar, "Approximation algorithms via contraction decomposition," in *SODA*, 2007, pp. 278–287.
- [16] E. Dinits, A. Karzanov, and M. Lomonosov, "On the structure of a family of minimal weighted cuts in graphs," *Studies in Discrete Mathematics (in Russian)*, ed. A.A. Fridman, 290–306, Nauka (Moskva), 1976.
- [17] J. Edmonds and E. Johnson, "Matching: A well-solved class of integer linear programs," in *Proceedings of the Calgary International Conference on Combinatorial Structures and Their Applications*, e. R. Guy et al., Ed., 1970, pp. 82–92.
- [18] T. Feder and M. Mihail, "Balanced matroids," in *STOC*. ACM, 1992, pp. 26–38.
- [19] T. Fleiner and A. Frank, "A quick proof for the cactus representation of mincuts," Egerváry Research Group, Budapest, Tech. Rep. QP-2009-03, 2009.
- [20] D. Gamarnik, M. Lewenstein, and M. Sviridenko, "An improved upper bound for the tsp in cubic 3-edge-connected graphs," *Operations Research Letters*, vol. 33, no. 5, pp. 467–474, 2005.
- [21] M. X. Goemans, "Minimum bounded degree spanning trees," in *FOCS*, 2006, pp. 273–282.
- [22] M. Grigni, E. Koutsoupias, and C. Papadimitriou, "An approximation scheme for planar graph tsp," in *FOCS*, 1995, p. 640.
- [23] M. Held and R. Karp, "The traveling salesman problem and minimum spanning trees," *Operations Research*, vol. 18, pp. 1138–1162, 1970.
- [24] P. N. Klein, "A linear-time approximation scheme for planar weighted tsp," in *FOCS*, 2005, pp. 647–657.
- [25] R. Lyons and Y. Peres, *Probability on Trees and Networks*. Cambridge University Press, 2010.
- [26] T. Mömke and O. Svensson, "Approximating graphic tsp by matchings," in *to appear in FOCS*, 2011.
- [27] C. L. Monma, B. S. Munson, and W. R. Pulleyblank, "Minimum-weight two-connected spanning networks," *Math. Program.*, vol. 46, pp. 153–171, 1990.
- [28] S. Oveis Gharan, A. Saberi, and M. Singh, "A randomized rounding approach to the traveling salesman problem," *available on arXiv.org*, 2010.
- [29] C. H. Papadimitriou and S. Vempala, "On the approximability of the traveling salesman problem," *Combinatorica*, vol. 26, no. 1, pp. 101–120, 2006.
- [30] C. H. Papadimitriou and M. Yannakakis, "The traveling salesman problem with distances one and two," *Mathematics of Operations Research*, vol. 18, pp. 1–11, February 1993.
- [31] R. Pemantle, "Towards a theory of negative dependence," *J. Math. Phys.*, vol. 41, pp. 1371–1390, 1999.
- [32] J. Pitman, "Probabilistic bounds on the coefficients of polynomials with only real zeros," *J. Comb. Theory Ser. A*, vol. 77, pp. 279–303, February 1997.
- [33] A. Schrijver, *Combinatorial Optimization*, ser. Algorithms and Combinatorics. Springer, 2003, vol. 2.
- [34] L. A. Wolsey, "Heuristic analysis, linear programming and branch and bound," in *Combinatorial Optimization II*, ser. Mathematical Programming Studies, 1980, vol. 13, pp. 121–134.